

Understanding Assurance Cases: An Educational Presentation in Five Parts

Module 2: Application

C. Michael Holloway
c.michael.holloway@nasa.gov

Senior Research Computer Engineer
Safety-Critical Avionics Systems Branch
NASA Langley Research Center, Hampton, Virginia, U.S.A.

UNDERSTANDING ASSURANCE CASES

MODULE 2: APPLICATION

C. MICHAEL HOLLOWAY

NASA LANGLEY RESEARCH CENTER
C.MICHAEL.HOLLOWAY@NASA.GOV

The past is never dead. It's not even past. - William Faulkner

VERSION 1.0A

This material was created in 2015-16, as part of the Explicate '78 project. The project was supported in substantial part by the Assurance Case Applicability to Digital Systems task under the reimbursable interagency agreement with the Federal Aviation Administration for Design, Verification, and Validation of Advanced Digital Airborne Systems Technology (IAI-1073 Annex 2 for NASA; DTFACT-10-X0008, Modification 0004 for the FAA). The original presentations were delivered to a selected group of FAA civil servants and NASA Langley personnel. The audio was recorded and partial transcripts (containing only the words spoken by the presenter, Mr. Holloway) produced. The intent from the beginning was to collect the material into a form that could be made available publicly. The text adheres closely to the original transcript, with the exception of an occasional insertion of new information that arose since the original presentation dates. The full collection consists of six documents (including this one), which are available electronically through <https://shemesh.larc.nasa.gov/arg/uac.html>.

Hello everybody.

Welcome to the second module in our educational series about Understanding Assurance Cases. In this module, we will examine the **Application** of the assurance case concept.

We'll be talking about the past and the present, hence a famous quote from William Faulkner encapsulates our theme:

“The past is never dead. It's not even past.”

[Faulkner, William. 1951. *Requiem for a Nun*. act i, scene iii. New York: Random House.]

As with Module 1, there will be several times when I'll stop to give you a chance to ask questions; but feel free to interrupt me at *any* point if you have a burning question. I'll either try to answer it right away, or defer it to a better time a bit later on.

Before going any further, I will repeat verbatim some preliminary remarks I made at the beginning of Module 1.

Within the assurance case community, intramural debates abound about a variety of topics we will discuss. Except in rare instances the existence of these debates is intentionally ignored or mentioned only briefly in this material. Here's why.

Disagreements exist about terms, definitions, notations, philosophy, procedures, tools, and just about everything else.

The depth of the disagreements ranges all the way from *shallow* differences in preferences (which term best denotes a particular concept, for example), to rather *deep* philosophical differences (the feasibility and desirability of formalizing assurance arguments, for example).

Spending *too much* time on these disagreements would likely make this material deeply confusing; but spending *too little* time on them might hinder your understanding of some materials you may come across.

In trying to strike a balance, what I've chosen to do is *not* highlight the areas of disagreement on the slides (except occasionally where it seems essential), but to mention the disagreements where appropriate in my words accompanying the slides.

One other quick note before we proceed: All images you see were either created by me (Michael Holloway) or are in the public domain via CCo 1.0 Universal. For images that do not fall into either category, you will see only links, not the actual image that was used in the original presentation.

Here are the four learning objectives for Module 2.

LEARNING OBJECTIVES

A person completing Module 2 should be able to

- ❖ Cite selected past events relevant to the development of the assurance case approach
- ❖ List uses of assurance cases in several domains
- ❖ Discuss possible lessons learned from past uses
- ❖ Explain potential benefits and problems associated with assurance cases

The past is never dead. It's not even past. - William Faulkner

By the time we're finished today, I hope that you'll be able to do at least these four things:

One, cite selected past events relevant to the development of the assurance case approach.


Two, list uses of assurance cases in several domains.

Three: discuss possible lessons learned from past uses.

And four: explain potential benefits and potential problems associated with assurance cases.

[Question to participants: Any questions about these learning objectives?]

As I hope you remember, in Module 1 we introduced Jon (the young fellow on the left) Jon's dad (named Mike, the not-nearly-as young fellow on the right), and Tim (a fellow we've not seen, but with whom Jon wants to ride to a game).



You also, I hope, recall that at the end of our story in Module 1 Jon's dad had said that he wanted to see an assurance case for why he should believe that Tim would get Jon to and from the game in one piece.

To this, Jon had replied, "I'll ask Tim if he has one."

A few hours later, the story continues as Jon and his dad get together again.

Jon says, "I asked Tim about an assurance case."

Jon's dad asks expectantly, "What did he have to say for himself?"

"He was a tad bit confused," replies Jon.

Mike, himself a tad bit confused, asks, "Confused? What was he confused about?"

Jon answers: "Whether an assurance case is the same thing as a safety case."

"Then he's heard about safety cases?" Jon's dad replies, once again expectantly.

Yes, from his sister
Rose, who lives in
England.



That's where it
all started,
you know.



“Yes, from his sister Rose who lives in England.”

“That’s where it all started you know.”

“Safety cases came first. Assurance cases are more general. All I really need from Tim is just a safety case.”

Jon breathes a sigh of relief, and exclaims, “That’ll make Tim happy. Thanks dad!”

After a short pause, Jon adds, with a slight smirk on his face, “Oh, Dad, I almost forgot ... Tim also wondered ... if you want a *brief case*, or a long one.”

Mike starts to reply, then the word play registers in his mind, and he simply smiles.

As Mike said, safety cases did come first. So, let’s talk a bit about the origins of safety cases. Like many origin stories, the details are a bit murky, and not everyone agrees about when and how things started.

Everyone *does agree*, that the beginnings were not so long ago in places not so far away from us. Back then and over there ...

Not so long ago in places not so far away ...

1957 Windscale – fire

1974 Flixborough – explosion

1976 Seveso – venting into atmosphere

1988 Piper Alpha – explosion; Clapham Junction – crash

1999 Ladbroke Grove – collision of local and high speed train

2006 Nimrod – loss of aircraft from an inflight fire after aerial refueling

... the accidents you see here were catalysts for changes in the way people thought about and regulated various dangerous activities and systems.

Windscale, Flixborough, Seveso, Piper Alpha, Clapham Junction, Ladbroke Grove, and, Nimrod have all played a part in safety case history and lore.

We could easily spend a whole hour or more talking about any one of these alone, so what I'll present are only incomplete overviews; and as with any overview, I may leave out some things that other folks would include, and include some things that other people would leave out.

Let us begin in northwestern England quite close to the coast, three and one quarter years before I was born

In October 1957, a fire at the Windscale nuclear reactor facility and plutonium-production plant resulted in a major release of radioactive materials. The fire started when a routine heating of the graphite control blocks in the number 1 reactor ran out of control, rupturing adjacent uranium cartridges. The uranium oxidized, causing a fire that burned for 16 hours before it was extinguished, and releasing radioactive Iodine-131 into the atmosphere, and also melting about 10 tons of the reactor core.

The UK government banned, for several weeks, the sale of milk from about a 200 square mile area around the site but generally told the public few details about the accident at the time. Windscale remains today the most serious nuclear power accident in the UK.

Not so long ago in places not so far away ...

1957 Windscale

Nuclear Installations Act →

Nuclear Installations Inspectorate →

Office for Nuclear Regulation <onr.org.uk>

A photograph of a helicopter flying over burning site is available at

<https://bbc.in/2PNIBcC>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

Among the actions taken in the wake of the accident was the adoption in 1959 of the Nuclear Installations Act

This act established the Nuclear Installations Inspectorate, which in turn required prospective reactor installations to submit a set of reports justifying the safety of the design, construction, and operation of the plant.

Although the term ‘safety case’ was not used in these early days, many, in retrospect, consider the certification process that was established in the wake of Windscale as the true beginning of the safety case approach.

Over the years, the UK commercial nuclear power regulations became increasingly more clearly safety-case based. Regulation is now the responsibility of the Office for Nuclear Regulation, which is an agency of the Health and Safety Executive. Their web site, O-N-R dot O-R-G dot U-K, is worth visiting, if for no other reason than to see in print an usually well-written mission statement: ‘The Office for Nuclear Regulation’s mission is to provide efficient and effective regulation of the nuclear industry, holding it to account on behalf of the public.’

They also have several short documents worth reading; we’ll come back to at least one of those in later Modules.

Let’s now move forward in time nearly 17 years, and in geography about 200 miles south east across England.

Not so long ago in places not so far away ...

1974 Flixborough

Advisory Committee on Major Hazards →
recommended draft regulations →
not enacted because superseded by events

A photograph of Flixborough five days after the event is available at

<https://bit.ly/2CWOAJf>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

To provide the summary of the Flixborough disaster, I'm going to quote directly from the report produced by the Court of Inquiry, because improving on its words, or at least its first sentence, is impossible.

“At about 4.53 pm on Saturday 1st June 1974 the Flixborough Works of Nypro (UK) Limited were virtually demolished by an explosion of warlike dimensions. Of those working on the site at the time, 28 were killed and 36 others suffered injuries. If the explosion had occurred on an ordinary working day, many more people would have been on the site, and the number of casualties would have been much greater. Outside the Works injuries and damage were widespread but no-one was killed. Fifty-three people were recorded as casualties by the ... police; hundreds more suffered relatively minor injuries which were not recorded. ... Property damage extended over a wide area.”

[Department of Employment. 1975. *The Flixborough disaster: Report of the Court of Inquiry*. London: Her Majesty's Stationery Office.]

Whether the explosion's initiating event was a failure in a 20-inch bypass or in an 8-inch pipe was the subject of much controversy during the court of inquiry and ever since. The inquiry came out in favor of the 20-inch hypothesis but the initiating event is not important for our purposes. What's important is that in response to the accident an Advisory Committee on Major Hazards was created within the Health and Safety Executive.

The Committee recommended that regulations be established to “ensure identification, assessment and management of potential hazards in chemical installations.” These

recommendations resulted in draft regulations, which were not enacted because other events outside of the UK happened to change the regulatory landscape.

Not so long ago in places not so far away ...

1976 Seveso

EEC “Seveso directive” → (EU “Seveso II”)
UK CIMAH regulations → (UK COMAH)
‘safety case’ term came to be used

A photograph of cleanup crews at Seveso is available at

<https://bit.ly/2S7awpv>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

These other events began on July 10, 1976, not in the UK, but rather in Seveso, Italy, a few miles north of Milan, when a rupture disc blew on a chemical reactor operated by the Icmesa chemical company. This occurred when a batch process was interrupted before the final step was completed (to conform to Italian law concerning hours that a plant could be operating). The interruption resulted in a spike in steam temperature, which was unseen by the operators because the vessel had no active temperature measurement.

The steam overheated the upper part of the reactor chamber, and with agitation turned off as part of the plant shutdown process, an exothermic runaway reaction began. This reaction produced tetra-chloro-di-ben-zo-p-di-oxin (known as TCDD, and sometimes incorrectly called simply dioxin), which is a highly toxic chemical that the plant *did not* produce during normal operations. No deaths were directly attributed to the TCDD release, but many people got sick, many animals died, and a substantial area had to be evacuated and stripped of soil.

The accident led to the European Economic Community adopting in 1982 what became known as the Seveso directive. The adoption of this directive is cited by some folks as the true origin of safety cases. Among other things the directive required member states to make manufacturers responsible for “tak[ing] all the measures necessary” to prevent “major accidents” and to “prove” that they have done so.

The United Kingdom implemented the directive through the Control of Industrial Major Accident Hazards (CIMAH) regulations in 1984. CIMAH required manufacturers to “provide evidence including documents to show that” they have ... “identified the major accident hazards; and ... taken adequate steps to ... prevent ... major accidents and to limit their consequences to persons and the environment, and ... provide persons working on the site with the information, training and equipment necessary to ensure their safety.”

Some other people consider the CIMAH regulations to be the true origin of safety cases, perhaps because the term itself came to be used in relation to documents produced by manufacturers to comply with the regulations.

The original Seveso directive has since been superseded by a European Union law generally known as Seveso II; in the UK the CIMAH regulations were replaced by the Control of Major Accident Hazard Regulations (COMAH). The safety case idea is still strong with them.

CIMAH applied to installations on-shore that posed major accident hazards. It did *not* apply to off-shore installations.

Not so long ago in places not so far away ...

1988 Piper Alpha

Lord Cullen inquiry →

“risks offshore are clearly no less” →

Offshore Installations (Safety Case) Regulations

Many photos of a burning Piper Alpha are available. One of them is here:

<https://bit.ly/2Cxl3VG>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

Piper Alpha was an off-shore installation (an oil platform to be specific) located in the North Sea about one hundred ten miles from Aberdeen, Scotland. On July 6, 1988, two hundred and twenty-six people were aboard the platform when it experienced a series of catastrophic explosions and fires. One hundred sixty-seven people were killed (including two not from the platform who died in a rescue attempt), and the platform was totally

destroyed. Because the platform was destroyed, little physical evidence was available for investigators, so the precise combination of events that led to the disaster is not known for sure.

The public inquiry led by Lord Cullen concluded that most likely the initial explosion occurred when a pump was restarted after maintenance by operators who were unaware that a relief valve in the pump discharge had also been removed for maintenance, and a blank loosely installed in its place. This blank leaked, producing a flammable hydrocarbon cloud, which found an ignition source. From that point things spiraled out of control in a variety of ways we won't take time to discuss now.

In addition to determining the likely direct causes of the accident, and discussing specific related recommendations, the inquiry by Lord Cullen also considered more general issues.

One resulting recommendation was that off-shore operations should be required to have a safety case just like on-shore operations. He wrote: "A Safety Case should be required for existing installations. This is the case onshore. The risks offshore are clearly no less. It is not acceptable that installations should be operated without a thorough assessment of what those risks are."

He further wrote that the Safety Case should be primarily "the means by which an operator demonstrated to itself the safety of its activities." Lord Cullen emphasized that the Safety Case should not be a static document, but part of a continuing dialog about safety, including between the operator and the regulatory body, whose role would largely be one of auditor.

As a direct result of Lord Cullen's recommendations, the Offshore Installations (Safety Case) Regulations were introduced in the UK in 1992, making the processing industries onshore and offshore subject to producing and maintaining safety cases.

The story turns now from processing to transportation, particularly rail transportation in the UK.

[Question to participants: Before I continue the story, does anyone have any questions?]

There are two pertinent rail accidents for us to discuss.

The first happened in 1988, the same year as the Piper Alpha disaster. I'll tell you about it by quoting some excerpts from report produced by the inquiry into the accident led by Anthony Hidden, because improving on its excellent wording is unlikely.

[Hidden, Anthony. 1989. *Investigation of the Clapham Junction Railway Accident*. London: Her Majesty's Stationery Office.]

"On the railway lines between Waterloo and Wimbledon four tracks run through a cutting a mile or so to the country side of Clapham Junction railway station. ... Just after 8 a.m. on Monday, 12 December 1988 three specific trains were running towards that cutting on their normal timetables. Two passenger trains were heading into Waterloo One, the 07:18 from Basingstoke, the other, running behind it from the South Coast, the 06:14 'Poole' train. The third train, the 08:03 Waterloo to Haslemere, was running without passengers ... on [an] adjoining line."

Not so long ago in places not so far away ...

1988 Clapham Junction

Anthony Hidden inquiry →
found railway safety culture wanting →
Railway (Safety Case) Regulations

A photograph of Clapham Junction post-accident operations is available at

<https://bit.ly/2PSZOS8>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

“At about 8:10 ... the driver of the ‘Poole’ train, having come into the cutting on his way into Waterloo ... and having passed signals in his favour at all stages, cleared the visual obstruction of the steep bank on the left-hand curve. At that moment he must have come upon what was, in signaling and therefore in driving terms, unthinkable and impossible: immediately ahead of him was the Basingstoke train on the same line, stationary, and within a distance in which the ‘Poole’ train could not possibly be stopped.”

“Despite full emergency braking of the ‘Poole’ train, its leading coach collided head-on with the rear of the Basingstoke train. The collision forced it out to its off-side where it struck the third ‘empty’ train going in the opposite direction. ... An appalling accident had happened.”

Thirty-five people died as a result of the accident (Thirty-three on scene, and two a bit later from their injuries. All of them had been carried in the first two coaches of the ‘Poole’ train.

The physical cause of the accident was fairly straightforward to uncover: a signal failure, which had been caused by a maintenance-induced wiring fault.

The inquiry, however, did not stop at finding the physical cause, it also discussed the whole railway safety culture at the time, and found it wanting. The inquiry’s report was one of the catalysts for a wider public consideration of railway safety, which ultimately led to the introduction in 1994 of Railway (Safety Case) Regulations, which required

railway infrastructure controllers and all train and station operators to prepare safety cases that demonstrated sufficient thought about and management of all credible hazards.

The Clapham Junction accident led to the requirement for safety cases in the railways; another accident more than a decade later led to deeper consideration of the content and disposition of such cases.

Not so long ago in places not so far away ...

1999 Ladbroke Grove

Lord Cullen inquiry:

“... poor quality of certain safety cases ...”

“... the application of the safety case ... is endorsed.”

A photograph of the derailed trains at Ladbroke Grove is available at

<https://bit.ly/2ytiKA5>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

“On 5 October 1999 at Ladbroke Grove junction, about two miles west of Paddington Station, London, there was a head on crash at high speed between trains operated by Thames Trains and First Great Western (FGW). This caused the death of [] 31 persons ... include[ing] both train drivers, and inflicted injuries, some of them critical, on over 400 other persons.”

Lord Cullen, of Piper Alpha fame, conducted a public inquiry into the accident and eventually published a 2-volume report. The first volume of the report dealt mainly with specifics of the accident. The brief summary I gave a moment ago comes directly from words in volume 1.

[The Rt Hon Lord Cullen, PC. 2000. *The Ladbroke Grove Rail Inquiry: Part 1 Report*. Norwich: HSE Books.]

The inquiry discovered that the Thames Train passed a Red danger signal travelling at about 41 mph, leading it to the Main line, on which the First Great Western high speed train was approaching on green signals. Both train drivers applied their brakes, but too

late to have any significant effect. The collision occurred at a combined speed of about 130 mph. The inquiry considered it more probable than not that the poor sighting of the signal passed at danger, coupled with bright sunlight at a low angle, were factors that led the train driver to think that he had a proceed aspect.

The second volume produced by Lord Cullen's inquiry "was concerned in regard to the railways, with the management of safety and the regulatory regime." Lord Cullen noted "The general object of a safety case is to ensure that an operator has the will, capabilities, organisation, system and resources to operate safely."

[The Rt Hon Lord Cullen, PC. 2001. *The Ladbroke Grove Rail Inquiry: Part 2 Report*. Norwich: HSE Books.]

He further stated: "The application of the safety case to Great Britain's railways is endorsed. ... there is a need for the framework provided by the Safety Case Regulations, within which the duty holder demonstrates, and by reference to which it operates, its arrangements and procedures for the management of safety in a consistent and effective manner."

Lord Cullen also noted "The Inquiry heard evidence from a number of witnesses about the poor quality of certain safety cases, especially the earliest which had been produced."

In discussing poor quality safety cases, he stated "While it is clear that the safety case can become overbureaucratic, it has the potential to be a valuable tool, by, for example, bringing about a systematic approach to safety and providing a record of management's commitments to safety. The evidence showed that it can be a 'living document', part of the direct management of safety. The discipline of producing a safety case has an important value in itself. ... The evidence [also] demonstrated the significance of ensuring employees' understanding and knowledge of its substance."

Thus far, we've talked about nuclear power, chemical processing of various sorts, and railways.

[Question to participants: Any questions or comments at before we continue?]

We now turn to the air.

The last specific accident I'll discuss happened over southern Afghanistan on September 2nd 2006. While on a routine mission in support of NATO and Afghani ground forces, RAF Nimrod X V 230 suffered a catastrophic mid-air fire, leading to the total loss of the aircraft and the death of all twelve on board.

I'll describe what happened borrowing liberally from the Nimod Review report, written by Queens' Counsel, now Sir, Charles Haddon-Cave.

[Haddon-Cave, Charles. 2009. *The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006*. London: The Stationery Office.]

Not so long ago in places not so far away ...

2006 Nimrod

RAF Board of Inquiry: significant errors in safety case →

Hadden-Cave independent review →

Safety case approach is **not** a panacea

A photograph of a Nimrod aircraft is available at

<https://bit.ly/2P9OwvF>

The link is to a web site outside of NASA as a service to users. The presence of a link is not a NASA endorsement of the site, nor is NASA responsible for the information collection practices of non-NASA sites.

“XV230 had taken off ... at 09:13 hours that morning. All went according to plan until [about 2 hours later] when, some 1½ minutes after completion of Air-to-Air Refuelling ..., the crew were alerted that something was amiss by two almost simultaneous warnings: a fire warning in the bomb bay and a smoke/hydraulic mist warning in the elevator bay.” ... the camera operator reported ‘we have flames coming from the rear of the engines on the starboard side’. ... the crew immediately commenced emergency drills and ... transmitted a MAYDAY whilst diverting to Kandahar airfield.”

“Faced with a life-threatening emergency, every member of the crew acted with calmness, bravery and professionalism, and in accordance with their training. They had no chance, however, of controlling the fire.”

The aircraft eventually exploded in the air.

An RAF Board of Inquiry presented findings about the causes of the accident in 2007. It concluded that either fuel overflowed from a blow-off valve during the refueling or (less likely) fuel leaked from a coupling or pipe; the fuel came into contact with an exposed element of the aircraft’s Cross-Feed/Supplementary Cooling Pack duct.

It also found that the Safety Case prepared for the Nimrod between 2002 and 2005 contained significant errors.

Shortly after the Board of Inquiry findings were made public The Secretary of State for Defence announced that an independent review would be conducted to look into the

broader issues surrounding the loss of the aircraft. Haddon-Cave was appointed to conduct the review.

His report was published on October 28, 2009, and has been frequently cited in a variety of contexts ever since. I was in London at a System Safety Conference when the report was made public.

Certain critics of safety cases are especially fond of quoting selectively from the report as evidence supporting their negative opinions. Haddon-Cave does indeed identify a number of problems that occurred with the Nimrod safety case: to put it bluntly it was rubbish.

To be a bit more specific, using Haddon-Cave's words for the most part ...

"[The] attitude to the [Nimrod Safety Case] was fundamentally affected by the prevailing malaise ... that, because the Nimrod had operated safely for over 30 years, it could be assumed that the Nimrod was 'safe anyway' and that, therefore, the [Nimrod Safety Case] exercise did not really matter." "[The contractor's] approach ... was flawed and undermined from the outset: it approached the task assuming 'safety' and viewed the [Nimrod Safety Case] task as essentially a documentary or paperwork exercise aimed at proving something that it already knew, i.e. that the Nimrod was safe."

Haddon-Cave noted that the primary purpose of "a 'Safety Case' is to 'identify, assess and mitigate' all potential significant hazards to pieces of equipment, platforms or installations, including hidden, or previously unidentified, hazards. ... the drawing up of a 'Safety Case' [is] merely a means to achieving this end, ... intended to provide a structure for critical analysis and thinking, or a framework to facilitate a thorough assessment and addressing of serious risks. Unfortunately, in the case of the [Nimrod Safety Case], the production of a 'Report' became an end in itself. Critical analysis descended into a paperwork exercise. "

So the real lesson taught by the tragic Nimrod accident is *not* (as some critics would have you to believe) that a safety case approach is a bad idea, but *rather* that a safety case approach is **not** a panacea. Creating a document that is called a safety (or an assurance) case does not mean that a good case has been made.

There's a lot more that could be said about the Nimrod Review, and about everything else I've mentioned so far, and that are lots of other things that I could mention that I've not mentioned at all, such as the beginnings of research groups at places such as the University of York and City University London, but we'll stop with the history at this point. I'll have more to say about research groups in Module 5.

[Question to participants: Does anyone have any questions about the history?]

Let's move now to talk a bit about current practice.

Discussing current practice is complicated by the paucity of publicly accessible, detailed information about existing industrial cases; Such cases are typically regarded as proprietary information, and thus not available to view. It seems fair to say, however, that the use of safety / assurance cases in real life can be roughly divided in four categories, which I'll now show you.

SAFETY CASES IN CURRENT PRACTICE

Fully established

- UK nuclear
- EU, Australian, NZ process industries

Recently established

- UK+ rail
- UK air traffic management and defence

Being established

- Global automotive
- US medical devices

Being explored

- US process industries
- Navy UAS

US CSB (2014). *Regulatory Report: Chevron Richmond Refinery Pipe Rupture and Fire*, No. 2002-03-I-CA.

Among domains in which (safety) cases are fully established are the UK nuclear industry (as I mentioned already in the discussion of Windscale), many of the EU process industries (think Seveso), and also process industries in Australia and New Zealand.

Recently established domains include rail in the UK (and much of the EU), UK air traffic management, and various aspects of UK defence.

Domains in which the use of cases is in the process of being established include the global automotive industry, and certain aspects of US medical devices, particularly infusion pumps.

Finally, domains that are exploring use include Some US process industries, and the US Navy, at least in respect to UAS.

The US Chemical Safety and Hazard Investigation Board published in 2014 what they call a “Regulatory Report” concerning the 2012 Chevron Richmond Refinery pipe rupture and fire. Nearly all of the report deals with whether a “safety case regulatory regime” might be appropriate, reaching the conclusion that the CSB believes that adopting attributes of “more robust safety management regulatory regimes from around the world” “would greatly enhance existing federal and California process safety regulations.” The report is available at

https://www.csb.gov/assets/1/20/chevron_regulatory_report_06272014.pdf

I believe it is important to note that the majority of existing experience in using safety cases has tended to involve ‘services’ rather than ‘systems’. It’s been more about how a plant is operated than about specifics of the design of a particular system within the plant.

Please don’t look at this slide as a definitive, all-inclusive breakdown of current practice; it is simply a rough breakdown, which I believe to be mostly accurate at the current time. Some other folks may dispute the categorizations in some areas, and may have additions, also.

One could certainly suggest that much more could be included in the “being explored” category (FAA and NASA, for example), but I’ve tried to restrict this listing to domains in which there exists evidence of active, real-life, practical activity of some sort, and not just research efforts. Concerning research efforts, we at NASA Langley published earlier this year (that is, 2015) a contractor report developed by folks from Saab Sensis and Dependable Computing that, among other things, discusses the results of a literature search looking for examples of published assurance cases.

[Rinehart, David J., Knight, John C., Rowanhill, Jonathan. 2015. *Current Practices in Constructing and Evaluating Assurance Cases with Applications to Aviation*. NASA CR-2015-218678.]

Here is an excerpt from a table in the report. I won’t go into details, but I will note the column that mentions some of the relevant standards or regulations that exist in certain domains.

Some examples

Name	Domain	Organizations*	Standards / Regulations*
Offshore Oil and Gas	Energy	U.K. HSE Norway PSA U.S. API & COS	U.K. SI 2005 No. 3117 API RP 75
GDA of Nuclear Plants	Energy	U.K. ONR & EA	ONR-GDA-GD-001
CAP 670 & 760	Aviation Infrastructure	U.K. CAA	CAP 670 CAP 760
WAM Preliminary Safety Case	Aviation Infrastructure	Eurocontrol	WAM PSC
Risk-Informed Safety Case	Aerospace Vehicles	NASA Office of Safety and Mission Assurance	NASA System Safety Handbook Vol. 1
Triton UAS	Aerospace Vehicles	U.S. Navy	NAVAIR INST 13034.4
RAF Nimrod	Aerospace Vehicles	U.K. RAF	U.K. MoD JSP318B U.K. Defence Std 00-56 U.K. MoD BP1201
European Rail SMS	Railways	European Railway Agency	E.U. Directives 2001/14/EC, 2004/49/EC, 2008/57/EC
U.K. Rail Safety Cases	Railways	U.K. HSE	(not known / obsolete)
ISO 26262	Automobiles	ISO	ISO 26262
Infusion Pumps	Medical Devices	U.S. FDA	FDA 510(k)
Generic Pacemaker Assurance Case	Medical Devices	University of Pennsylvania	(none)

Rinehart, David J.; Knight, John C.; Rowanhill, Jonathan. Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation. NASA CR-2015-218678. January 2015.

In particular, there seems to be some evidence that the voluntary automotive standard ISO 26262 is growing in influence¹, which is a major reason that I included the global automotive industry in the *being established* category.

Concerning the FDA, based on conversations that I and others have had with some folks within that organization, they seem to be experiencing some of the same things that were highlighted in the Ladbroke Grove and Nimrod discussions earlier: namely, that some assurance cases are quite badly done.

I'm ready now to move on to discuss matters directly related to our 3rd and 4th learning objectives for this module, but before I do so, does anyone have a question?

Concerning possible lessons taught from the past, I think that a helpful way to think about such lessons is by using the famous five Ws, because history and practice seems to suggest that the answers to the questions “Who? What? Where? When? Why?” matter a great deal when it comes to applying assurance cases.

THE 5Ws MATTER A GREAT DEAL

Who?

What?

Where?

When?

Why?

In fact, they matter so much that one might perhaps accurately say that the overarching lesson taught thus far is “Carefully ask and answer the 5Ws.”

¹ Since this module was developed in 2015, the influence of the standard has continued to grow.

Let's look at each of these questions briefly.

THE 5Ws MATTER A GREAT DEAL

Who?

- Duty holder
- Developer
- Writer
- Operator
- Maintainer
- Regulator
- Assessor
- Public

To successfully apply an assurance case approach, you need to understand well *who* is necessarily involved in the application: eight categories of *whos* are listed here.

The *duty holder* refers to the people or organizations obliged to preserve the properties we're concerned about in the system or service (so, for example, safety for a safety case).

For simplicity, let's assume for the rest of this discussion that we're interested in safety of a system, so we don't have to say 'system or service' and we can talk about a 'safety case'. Let's further assume that the term 'organization' is a short hand for the longer phrase meaning something like "person, people, organization, or organizations."

The *developer* refers to the organization that will make design decisions and implement the system.

The *writer* refers to the organization that writes the safety case. One can argue that the Nimrod example suggests that it may not be a good idea to have an organization involved whose sole role is that of *writer*; though it is certainly true that the writing of safety cases has sometimes been contracted out to third parties without dire consequences.

The *operator* refers to the organization who will operate the system. In the process industries, the operator is often also the duty holder. This need not be the case. For example, consider pilots and drivers.

The *maintainer* refers to the organization who will service the system during its operational life. The maintainers might be disjoint with the duty holders, developers, and operators. For example, an operator might delegate maintenance to a contractor.

The *regulator* refers to the organization explicitly legally tasked with supervision of a particular system. Not all systems have regulators.

The *assessor* refers to the organization that explicitly audits or assesses the safety case. An assessor might or might not be a regulator. For example, an automobile manufacturer conforming to ISO 26262 might hire an outside organization to assess conformance even though law or regulation does not require conformance.

Finally, the *public* refers to anyone who may use or in any way be affected by the system in question and who is not in one of the other seven categories.

Those are the eight *whos* that are necessarily involved in the application of an assurance case approach. Trying to apply assurance cases without thinking about these *whos* is not a good idea.

[Question to participants: Are there any questions about “Who?”]

We’ll consider “What?” and “Where?” together.

THE 5Ws MATTER A GREAT DEAL

- Scope of design authority

- Scope of analysis

What?

- Intended

- operation

- environment

- safety target

Where?

- Scope of safety obligation

- Scope of system or service

These two questions encompass issues about scope and intent.

Let's talk about intent first.

What and where is the *intended operation* of the system? (Again we're using system as a shorthand for system or service). What's the system supposed to do? What's its mission? What is it *not* supposed to do? What are foreseeable types of misuse?

What and where is the *intended environment* of the system?

This includes any and all features of the (intended) time or place of operation. For example, will an aircraft perform taxi, takeoff, and landing in conditions of extreme cold, or extreme heat, or sandstorms?

What is the *intended safety (or assurance) target*?

This question might be answered By identifying an appropriate standard that will be followed (for example, software contributions to system risk should be controlled by applying DO-178C) or by identifying an appropriate risk acceptance test (for example, risks should be managed As Low as Reasonably Practicable - ALARP, So Far As Is Reasonably Practicable - SFAIRP, or Globally at least as good - GAMAB) As an aside: We could spend a lot of time talking about differences among these, but we won't (ALARP is hazard based. SFAIRP is precaution based. GAMAB is comparison based.)

"What?" and "Where?" also involve issues concerning scope.

Scope of design authority refers to what the developers are able to control.

Scope of analysis refers to what part of the system is being considered in the assurance case. Perhaps it is not the entire system but only certain aspects of it.

Scope of safety obligation refers to what the duty holder is obliged to consider.

Scope of system or service refers to the full extent of what we need to consider the safety implications of. This is very much related to intended operation.

Trying to apply assurance cases without thinking about these *whats* and *wheres* is not a good idea.

[Question to participants: Are there any questions about "What?" and "Where?"]

The "When" question concerns the timing of the creation of an assurance case or cases. Some possibilities are shown here on the slide.

THE 5Ws MATTER A GREAT DEAL

- Pre-operational case
- Operational case
- Maintenance case
- Subsystem / Subservice case

When?

A *Pre-operational case* comes from the system developers. The scope is limited to the system design and implementation, with operations assumed for the purpose of safety analysis. It is used to make release-to-service decisions. There can and should be several early versions of the pre-operational case. Early versions will describe the system as it will be (as far as is known at the time of writing); the final pre-operational case should describe the system as actually built. Note that the pre-operational cases necessarily lack evidence from experience of operation, and thus are based on assumptions about operation.

In an *operational case* the scope is the actual operation of the system in real life. System design issues are excluded (except for modification and monitoring). It addresses safety of operators and (if relevant) the public. It relies on a pre-operational safety case for claims about what the system does and (if appropriate) supports that case with information that shows that assumptions made in the pre-operational case about operations are correct.

Note that if the developer is not the operator the writer of the operational safety case might not be the same as the writer of the pre-operational case.

A *maintenance case* concerns, as you may suspect, how the system is being maintained, and should include discussion about the safety of the maintainers, and arguments concerning the maintenance assumptions made in the pre-operational case.

For some systems there may be separate cases for components or *subsystems*. These might be cited by the overall pre-operational, operational, or maintenance cases to

justify conclusions about the component or subsystem contribution to system risks or their mitigation.

In applying assurance cases the “When” question must be asked and answered. There’s some evidence to suggest that answering it with a single time may tend to be unwise. That is, writing a pre-operational case only, while ignoring operational and maintenance cases may fail to ensure the level of safety that is desired.

[Question to participants: Are there any questions about “When?”]

The final W question in this discussion, but almost certainly the first in a temporal sense, is “Why are you doing it?”

THE 5Ws MATTER A GREAT DEAL

- Communicate the rationale for believing that the system or service is acceptable for its intended use
- ⊗ Simply satisfy regulatory requirements

Why?

There’s a good answer and a bad answer to this question. Creating a case simply to satisfy regulatory requirements is the bad answer. Creating an assurance case to communicate the rationale for believing that the system or service is acceptable for its intended use is the good answer.

Each of the “Who” parties we talked about earlier should gain something from this communication. For example, Consider a pre-operational safety case written by the developer, who is also the duty holder.

The writer / developer / duty holder, who must articulate the rationale, might gain a more detailed understanding of that rationale, and recognize possible deficiencies.

A regulator might gain insight into whether applicable law or regulation has been complied with. An auditor might gain insight into what the duty holder considers adequately safe, what hazards they think are most in need of attention, what options were considered, and how they have gone about implementing the chosen options.

An operator might gain a better understanding of what a system or service is meant to be or do in order to be safe, thus putting that operator in a better position to notice operational realities that would make the system or service less safe than intended. A maintainer might gain a better understanding of which hazards a system's developers considered most in need of addressing and how they intended to address them.

Finally, if given access to the safety case, The public, whom might be harmed by the system or service, might gain a better understanding of how safe 'adequately safe' actually is.

This discussion leads us naturally into talking a bit about potential benefits which I've summarized here on this slide in four points.

POTENTIAL BENEFITS

- ❖ Improved, shared understanding amongst all relevant parties of hazards, vulnerabilities, ... , risks, and controls
- ❖ Greater focus on things that really matter
- ❖ Increased flexibility to use state-of-the-art methods, tools, approaches, ...
- ❖ More efficient and effective regulation

The first of these is directly related to what we've just discussed: An improved, shared understanding amongst all relevant parties of hazards, vulnerabilities, ... , risks, controls (and other things you might want to put here.)

There is also the potential for a greater focus on things that really matter, and for increased flexibility to use state-of-the-art methods, tools, approaches, and whatever else can be state-of-the-art.

Consequently, these things possibly could lead to more efficient and effective regulation.

These benefits are not givens, however, as we saw in several of the examples from history we discussed earlier.

POTENTIAL PROBLEMS - 1

❖ Cases can be used badly in many ways

- Failing to consider the 5Ws
- Relying on notation, automation, 3rd parties
- Failing to employ industry best-practices
- Treating the case as a product unto itself
- Failing to update the case when changes occur
- Listening to the wrong 'experts'
- Failing to pick the right level of detail
- ...

Cases can be used badly in many ways; I've listed seven of them on this slide.

Failing to consider the 5Ws.

Relying on notation, automation, or third parties.

Failing to employ industry best practices.

Treating the case as a product unto itself.

Failing to update the case when changes occur.

Listening to the wrong 'experts' (with the growing popularity of assurance cases, there's also a growing number of folks who style themselves as experts, but not all of them know what they're talking about).

Failing to pick the right level of detail.

Doing cases badly are not the only potential problems.

POTENTIAL PROBLEMS - 2

- ❖ Knowledge and skills may be required that are not abundantly present
 - within the developers
 - within the regulators
- ❖ Regulatory environment may not adequately empower the regulator to insist on good cases

Two others include the possibility that knowledge and skills may be required that are not abundantly present within the developers, or within the regulators; and, the possibility that the regulatory environment may not adequately empower the regulator to insist on good cases.

Even if a regulator has adequate skills, if the regulatory environment does not allow them to reject poor assurance cases, problems will certainly occur².

We're almost done, but before taking questions and comments, I want to show and read to you a superb quotation from the Haddon-Cave report.

² In the three years since this module was first presented, the importance of this particular problem relative to the other problems listed seems to have increased.

“At all stages of the safety pilgrimage it is vital to ask questions such as ‘What if?’, ‘Why?’, ‘Can you explain?’, ‘Can you show me?’, ‘Can you prove it?’. Questions are the antidote to assumptions, which so often incubate mistakes.”

“A Questioning Culture is the key to a true Safety Culture. In my view, people and organisations need constant reminding of the importance of asking questions rather than making assumptions, of probing and testing rather than assuming safety based on past success, of independent challenge of conventional wisdom ..., of the exercise of judgment rather than retreat behind the assignment of arbitrary quantitative values.”

“Questioning is a catalyst for thinking. As Professor McDermid told me, if he could replace all of the regulations with one word it would be: ‘THINK’”.

Haddon-Cave, C. (2009) *The Nimrod Review*. London: The Stationary Office. p. 574.
www.official-documents.gov.uk/document/hc0809/hc10/1025/1025.pdf

At all stages of the safety pilgrimage it is vital to ask questions such as “What if?”, “Why?”, “Can you explain?”, “Can you show me?”, “Can you prove it?”. Questions are the antidote to assumptions, which so often incubate mistakes.

A Questioning Culture is the key to a true Safety Culture. In my view, people and organisations need constant reminding of the importance of **asking questions** rather than making assumptions, of **probing and testing** rather than assuming safety based on past success, of **independent challenge** of conventional wisdom or shibboleths, of the **exercise of judgment** rather than retreat behind the assignment of arbitrary quantitative values.

Questioning is a catalyst for thinking. As Professor McDermid told me, if he could replace all of the regulations with one word it would be: “THINK”.

In my opinion the greatest potential benefit of the assurance case approach is that, used properly, it can force people to think more deeply than they otherwise would.

The greatest potential problem of the assurance case approach is that, if used improperly, it can cover up shoddy thinking.

[Question for participants: Any questions or comments before we end by reviewing the learning objectives?]

At the beginning, I listed four things that I hoped you’d be able to do by the end of this module.

Here are those four things recast in the form of questions. Think to yourself how you'd answer these questions.

REVIEW OF LEARNING OBJECTIVES

Are you able to

- ❖ Cite selected past events relevant to the development of the assurance case approach?
- ❖ List uses of assurance cases in several domains?
- ❖ Discuss possible lessons learned from past uses?
- ❖ Explain potential benefits and problems associated with assurance cases?

The past is never dead. It's not even past. - William Faulkner

After you've thought about the questions for a little bit, please ask me any questions that you still have for me.

If you have questions or comments about this material, contact its author at c.michael.holloway@nasa.gov.